Using Degradation Modeling to Identify Fragile Operational Conditions in Human- and Component-driven Resilience Assessment

Daniel Hulse
Intelligent Systems Division
NASA Ames Research Center
Moffett Field, CA, United States
Email: daniel.e.hulse@nasa.gov

Lukman Irshad
Intelligent Systems Division (KBR, Inc.)
NASA Ames Research Center
Moffett Field, CA, United States
Email: lukman.irshad@nasa.gov

***Abstract*—Studying failure events shows that many high-impact events result from the complex interactions between precipitating failure events and degraded operational conditions. Often, when a system is put in operations, unforeseen practical realities (e.g., maintenance and/or workforce availability) lead the system to be operated in configurations outside its envisioned nominal range. However, design-time failure models often assume that the failure events are initiated in an idealized, nominal state of system operation, resulting in an incomplete assessment of future risk. To solve this, this paper develops a framework to consider degraded operational performance in scenario-based resilience models which uses a corresponding model of performance degradation to determine the values of deteriorated model parameters in the resilience model. This framework is demonstrated on a remotely-piloted rover to determine the (individual and combined) effect of drive-train wear and operator fatigue on the resilience of the rover to drive-train faults. This demonstration showed the substantial impact that degradation has on resilience, highlighting the need to account for degradation in resilience models–specifically, unconsidered degradation can lead to overestimates of resilience (and thus underestimates of safety margin) and because resilience can degrade prior to visible unreliability, which can lead to an operational environment with a high propensity for high-impact unforeseen failure events.**

***Index Terms*—resilience, simulation, degradation**

I. Introduction

In 2009, Air France flight 447 from Rio de Janeiro to Paris crashed, resulting in 228 fatalities [1]. This accident began when severe weather system resulted in a temporary malfunction in the pitot probes, leading to the loss of airspeed measurements. Failing to grasp this, the crew, which at the time was being led by the least experienced of the three pilots [2], applied a series of maneuvers that caused the aircraft to stall and crash [3]. This accident resulted from degradation–first of pitot tubes (due to weather) and then of human performance (due reduced operator experience, mental workload, and stress), which reduced the crew's ability to identify and correct for the hazard. Preventing accidents like the Air France flight 447 thus requires understanding the effect of operational degradation on both components and human operators so that it can be mitigated before it causes an unsafe condition.

One of the most important developments in the mitigation of risk in design has been the shift from "random" failure models towards understanding how failures arise from degradation in system condition. This shift first happened in the field of reliability engineering, where component degradation (otherwise known as the "physics of failure") resulting from physical system failure mechanisms and external conditions came to be used to determine the failure rate of the system over time and resulting life-cycle reliability [4]–[6]. This approach enables improved design because understanding these mechanisms motivates one to think of the failures in terms of the initiating mechanisms (i.e., the root causes), resulting in systems which are reliable by-design, resulting in less maintenance and downtime [7], [8]. It further enabled the continuous prediction of reliability and maintenance, sparking the field of Prognostics and Health Management (PHM) [9], [10]. In PHM approaches, degradation models are used to determine the remaining useful life of the system based on when the changing failure rate (or some bell-weather health state) crosses a threshold. This approach enables the reduction of maintenance and increase in reliability by only taking the system out of service when needed to (rather than via a specified maintenance schedule).

While degradation was a marked methodological improvement over previous "random" failure models [9], [11], there are still many challenges to applying them in the design of new systems: (1) different components have different failure mechanisms, (2) data about these new components may be non-existent, (3) it requires the designer to have identified all the relevant mechanisms and risks, and (4) the general approach neglects failures resulting from manufacturing processes which cause early burn-in failures [8], [12]. A variety of techniques have further been used for degradation modelling, including monte carlo methods [11], bayesian networks [13], petri nets [14], neural networks [15] and others. Creating (physics-based, data-based, or hybrid) degradation models is

This research was funded by the System-Wide Safety project in the NASA Aeronautics Research Mission Directorate. The findings herein represent the research of the authors and do not necessarily the view of the federal government or NASA.

thus a key part of developing PHM technology [16], which has resulted in the development of degradation-oriented simulation tools for prognostics [17]. While these approaches generally involve modelling the system degradation, they do not consider the resulting performance and resilience effects–instead focusing on when to take the system out of commission to avoid unreliability.

However, degradation can have effects on system performance aside from its direct effect on component unreliability, such as lost performance and efficiency [18], [19]. This is particularly salient in gas turbines [18], [19] and human factors modelling [20], where it may not be economical to take the system out of commission to maintain nominal performance. Since these (and the previous) models show that performance, reliability, and remaining useful life degrades over time, it follows that resilience should also degrade due to the resulting decreased system capacity and increased fragility. However, very little has been done to take this into account in design-time resilience models. State-of-the-art modelling frameworks generally assume that there has been no deterioration in the system prior to the hazardous events [21], while frameworks which do account for deterioration consider it as change in fault probability [22]–which effects the likelihood of scenarios considered in a resilience assessment, but not the modelled recovery/resilience. In general, resilience metrics which approach "deterioration" or "degradation" consider it to be after the hazardous occurs [23], rather than a precursor which effects the resulting response. Thus, there is a need in the resilience field to better consider the deterioration in system state which happens prior to a fault event to better represent the true resilience of the system in operations.

A. Contributions

Thus, the aim of this work is to predict the loss of resilience due to system degradation to better understand the true resilience of a system in operations. To address this aim, this work creates a simulation framework to analyze degradations which originate from both humans and components acting across multiple time-scales. This framework, described in Section II involves using a model of system performance degradation to determine the parameters of a system resilience model for use in nominal and faulty simulations. This framework is then demonstrated in Section III, in analysis of a rover drive system to human (experience, stress, and fatigue) and component (wear) degradations, demonstrating how this approach can be applied over multiple time-scales.

II. Degradation-Based Resilience Assessment

The framework presented in this research supplements the system failure model with a corresponding degradation model that simulates performance-affecting states of the system over its lifecycle. This can be used to model both component and human-related degradations, where the degraded states of interest for the humans are performance shaping factors (e.g., operator experience, shift change-related changes to performance shaping factors, etc.). The degradation model is simulated over time and then sampled at given times to determine the degraded parameters of the system. These parameters are then simulated in the hazard model without introduced failure events to determine the degraded performance. Then, to determine how the system responds to hazards, failure events are injected and simulated in the hazard model, resulting in the degraded resilience of the system.

A. Degradation Modelling

Degradation models are used to determine the deterioration of system performance and reliability arising from loading, external conditions, and damaging physical processes over system's life-cycle. This deterioration, illustrated in the left (purple) section of Fig. 1, is a dynamic property which may vary depending on stochastic internal and external parameters, such as loading and process uncertainty, and may contribute to early failure of the system if too much deterioration occurs. A degraded state of the system D_1 can thus be modelled as the result of a dynamic process $d(t)$ over the system's operational life t_l:

$$D_1 = \int_{t \in t_l} d(t)dt \tag{1}$$

This represents a given system degradation provided the dynamics are deterministic. However, degradation processes often have stochastic behaviors in them (e.g., usage), making it subject to some behavioral uncertainty. The expected degradation D over a number of stochastic states O is thus:

$$D = \mathop{\mathbb{E}}_{o \in O} \{ \int_{t \in t_l} d(t,o)dt \} \tag{2}$$

B. Resilience Assessment

Resilience modelling is used to determine the system's response (and thus recovery) to hazardous scenarios. While a number of formalisms exist for resilience modelling, including network models [24], Bayesian graphs [25], and other timeless/single-timestep models, this work builds on the popular dynamic simulation approach. In this approach, the hazardous event is instantiated in a model of the system dynamics to determine the system's post-event behavior. This post-event behavior corresponds (shown in the right/orange section of Fig. 1) to the resilience triangle [26] used commonly in the literature [27], [28], showing the post-event degradation, robustness, and recovery. Considering a single hazardous event, the corresponding resilience state R_1 is thus the integral of the dynamic post-event behavior $r(t)$ over the interval t_1:

$$R_1 = \int_{t \in t_1} r(t)dt \tag{3}$$

Previous work has developed the fmdtools simulation toolkit [29] to conduct this sort of assessment, which further extends the resilience simulation concept from the system's response to a single event to its expected response to a set of hazardous scenarios. Using this definition, the resilience R is thus the expectation of the resilience state R_s over the set of hazardous scenarios S:

$$R = \mathop{\mathbb{E}}_{s \in S} \{ \int_{t \in t_s} r(t,s)dt \} \tag{4}$$

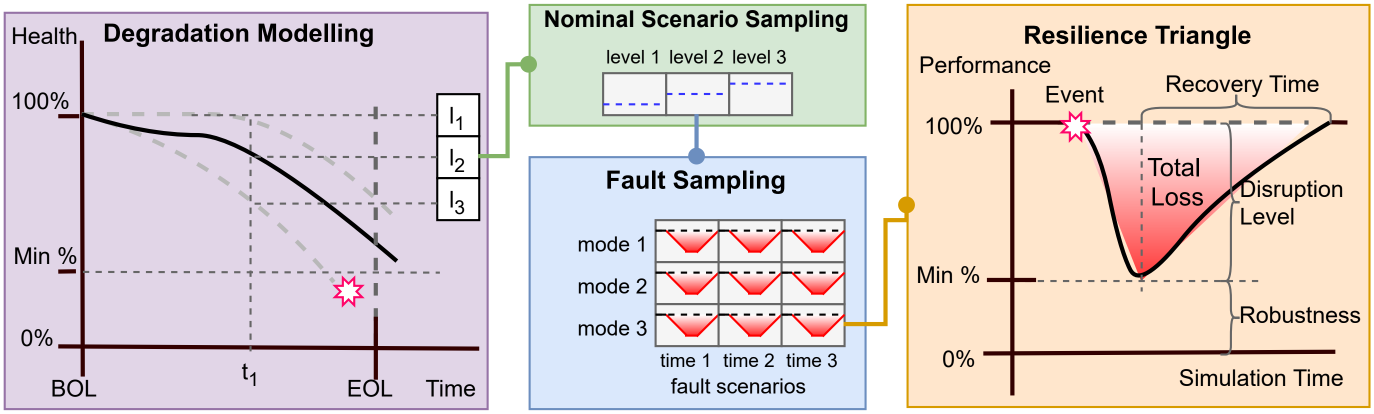

Fig. 1: Combined Degradation and Resilience Assessment Approach

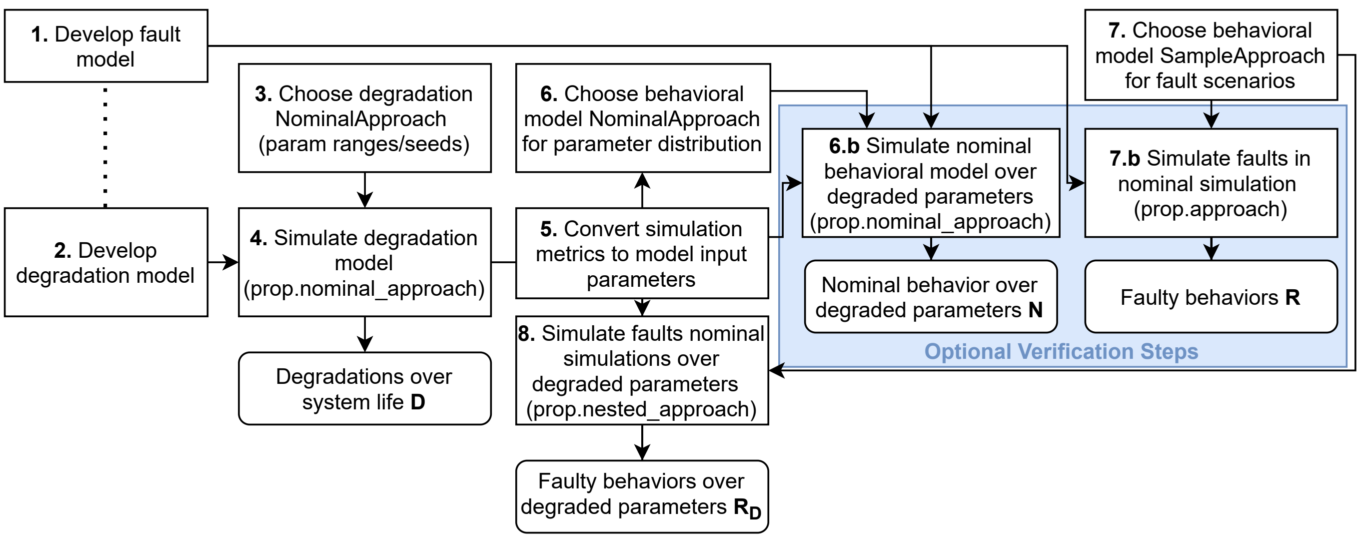

Fig. 2: Process for modelling fault resilience over set of performance degradations in fmdtools.

While this prior definition extended the concept of resilience to an expectation over multiple scenarios, the sampling approaches it used did not account for degradation in system performance prior to the scenario, instead solely considering the sources of hazards to be fault modes occurring at different times in the simulation interval. However, there is ongoing work to jointly consider external and internally-driven hazardous events by using a scenario sampling approach in which internal fault modes are simulated in a model which was instantiated at varying (nominal or off-nominal) parameters [30], as shown in the central (green and blue) section of Fig. 1.

C. Combined Concept

In this work, we further extend the nested scenario sampling approach by modelling operational parameters to be a result of a degradation model, creating the full framework shown in Fig. 1. This framework enables the consideration of operational deterioration in resilience assessment by nesting the simulation of faults in the simulation within the sampling of parameters resulting from different levels (and scenarios) simulated in a degradation model. Considering the expectation-based definition of resilience used in Eq. 4 and the similar definition of degradation in Eq. 2, the resulting degraded resilience R_D is:

$$R_D = \mathop{\mathbb{E}}_{o \in O}\{\mathop{\mathbb{E}}_{s \in S}\{\int_{t \in t_s} r(D(o), t, s)dt\}\} \quad (5)$$

where $D(o)$ is the result of the degradation function for a particular operational scenario o.

Developing this combined assessment of degradation and resilience follows the process shown in 2. As shown, first the fault model and degradation models must be developed such that they can be simulated. It should be noted that while these models likely share some correspondence in terms of structure, they can be developed separately as long as the outputs of the degradation model correspond to input parameters for the fault model. The degradation model is then simulated over a set of scenarios determined by the designer, the results of which (D) are sampled over various times to determine the set of degraded parameters. Using these parameter values, the

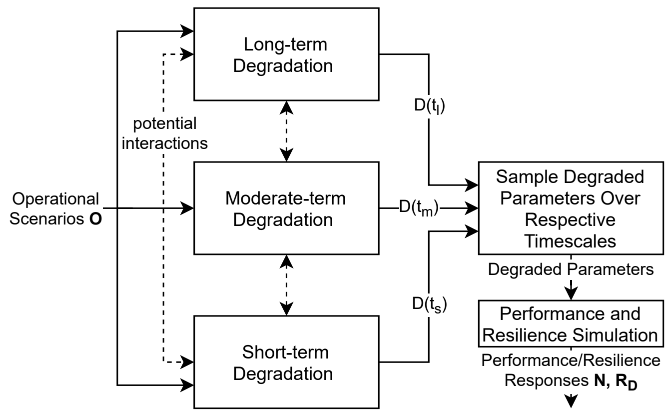

Fig. 3: Assessing Degradation that Occurs at Multiple Scales

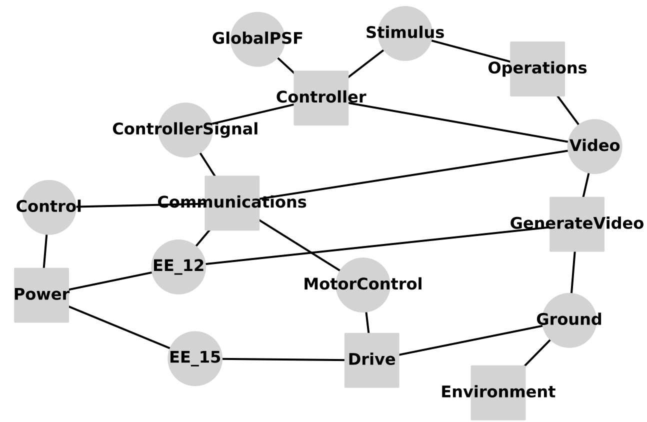

Fig. 4: Rover Model Functions and Flows

model can then be simulated in the nominal state to determine the effect on nominal performance, and additionally simulated over the set of fault scenarios to determine the degraded resilience at the given level/scenario of degradation.

D. Multi-Scale Degradation

Finally, it should be noted that degradations can occur across multiple time-scales and in different parts of the system. For example, while the state of health (i.e., capacity) of a battery may degrade over its operational life, the state of charge of the same battery may degrade over a day as a part of the routine usage and charging cycle. Since both conditions can simultaneously impact the available resilience of the system, it is important to not just these degraded behaviors/parameters individually, but joint together. To evaluate these joint degradations, we propose the framework shown in Figure 3. As shown, the degradation models are simulated over their respective time-scales and then sampled at the times of interest to determine the joint-degraded parameters D to use in the resilience simulation. Depending on the interactions between short and long-term degradation models (e.g., between long-term battery capacity and end-of-cycle charge), these may be simulated independently or in a nested approach, with the outputs of the long-term degradation model feeding the moderate and/or short-term models.

E. Human Performance Degradation

Using degradation models to account for human performance has specific challenges which must be considered during model setup. First, different human degradation behaviors may occur at different timescales, including the timescale of the resilience simulation itself. For example, mental workload may vary from mission to mission and influence the accumulation of fatigue in the short term. Accounting for human degradation in this situation would thus require a multi-scale degradation modelling framework presented in Section II-D in which the accumulation of fatigue results from different workloads over a day. In a situation like this, careful attention should be paid to ensuring that the combined degradation modelling framework does not lead to divergent loops (e.g., workload increasing fatigue and fatigue increasing workload), which could be accomplished by using a stand-alone model for mission degradation (rather than the resilience model itself) to feed the short-term degradation model. Second, models of human performance degradation need to be based on justified theory to avoid methodological errors, repetitions, and over/underestimations. For example, if one misses the inverse relationship between experience and stress, it may lead to over/underestimation of stress, resulting in a degradation model that is not realistic. Using existing human factors models to represent performance shaping factor dependencies (e.g., NUREG/CR-6883 [31]) and interactions (e.g., Refs. [32]–[34]) in cases like this can ensure that the degradation model results are accurate and theoretically-justified. Third, performance shaping factors can be influenced by a number of sub-factors which should be accounted for in the degradation model. For instance, experience can be influenced by training intensity, training frequency, and knowledge recalls [33]. Accounting for these influences in a degradation model could be achieved by modeling the experience accumulation rate as a function of these training sub-factors.

III. Rover Model Demonstration

To demonstrate this framework, this paper presents the design of a semi-autonomous rover. This rover was modelled at a high level to perform a basic autonomous driving task, with the functional model of the model shown in Fig. 4. As shown, this model encompasses the rover power and control systems, as well as its drive system, avionics, and interactions with its environment (i.e., movement and position with respect to a map). The task is to follow a given line from a given starting location to a given end location. While many different input lines can be used for different routes, the route used in this paper for demonstration purposes follows a simple sine curve as shown in Figure 5. If the rover deviates from the center line, it may go off course and crash into its surroundings. When the distance from the center line is greater than 1m, the rover can no longer see the center line and stops moving, because the rover has crashed. The main consideration used

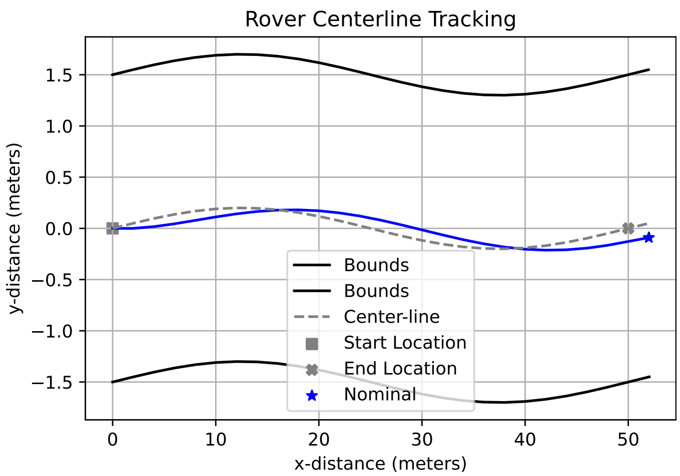

Fig. 5: Rover Model Environment

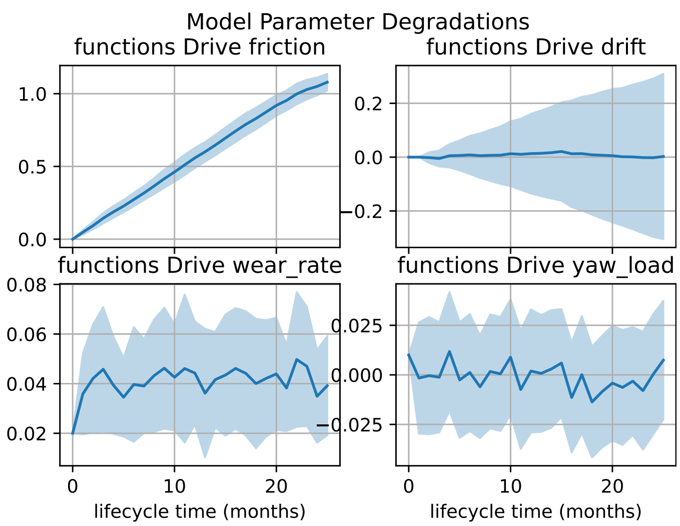

Fig. 6: Drive Degradation Model Results

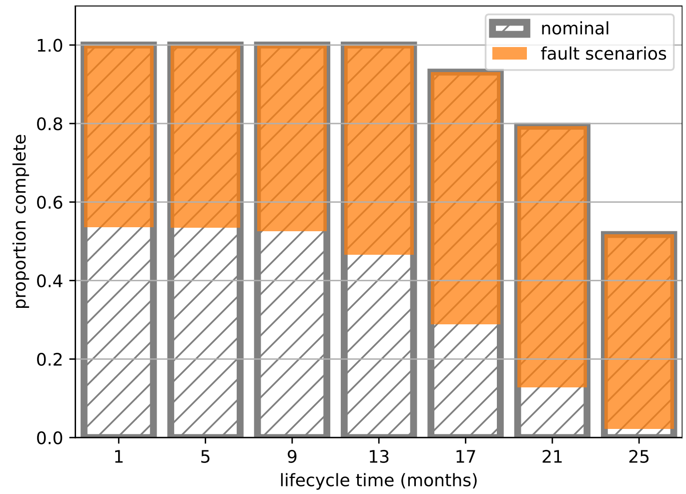

Fig. 7: Rover Resilience over Degraded Drive Parameters

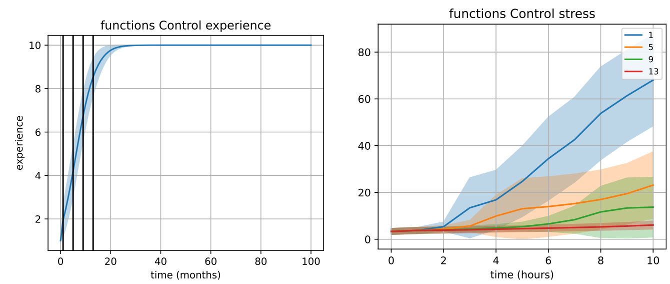

(a) Experience

(b) Stress

Fig. 8: Human Degradation Model Results

here for mission success/failure in this analysis is whether the rover makes it to the desired end-location. Additionally, for consideration of simplicity and computational costs, the resilience analysis focuses on faults in the drive-train (e.g., mechanical breakage, a stuck wheel, loss of powerplant, etc.).

A. Drive System Degradation Analysis

To show how component degradation models can be used to inform resilience assessment, this section studies the effect of drive degradation on the rover's ability to complete its mission. The main degradation behavior considered was the wear on the drivetrain, which results in more friction over time, and a drift in the drive-train direction from its intended course. This model is shown in Figure 6. As shown, these drive parameters degrade quickly over the notional 25 months because of the high wear rate and lack of maintenance. To evaluate the rover's performance and resilience given this degradation, the model was then simulated with friction and drift parameters corresponding to the outputs of the degradation model at multiple time-steps during the degradation.

The resulting effect of degradation on performance and resilience to drive faults is shown in Figure 7. As shown, the increased friction and drift over time leads the system to have reduced reliability at the 17^{th} month, resulting in a decreased proportion of nominal missions completed. At the 25^{th} month, only roughly 50% of missions are successfully completed, even if there are no fault scenarios. However, as shown, the resilience of the system also degrades considerably. While the system is able to mitigate $\tilde{5}4\%$ of the faulty scenarios at the first months, this proportion begins to degrade at the 13^{th} month–before reliability is effected. As a result, at the 25^{th} month, only 2% of fault scenarios in the drive system are able to be mitigated. This demonstrates how component degradation can have a large effect on resilience–because resilience often degrades prior to unreliability, it can lead to a false sense of security that the system is safer than it is.

B. Human Degradation Analysis

To further show how (1) this degradation modelling framework can be used to evaluate the resilience effects of degraded human performance and (2) demonstrate the multi-scale degradation approach presented in Section II-D, this section studies the how the interaction between operator experience (devel-

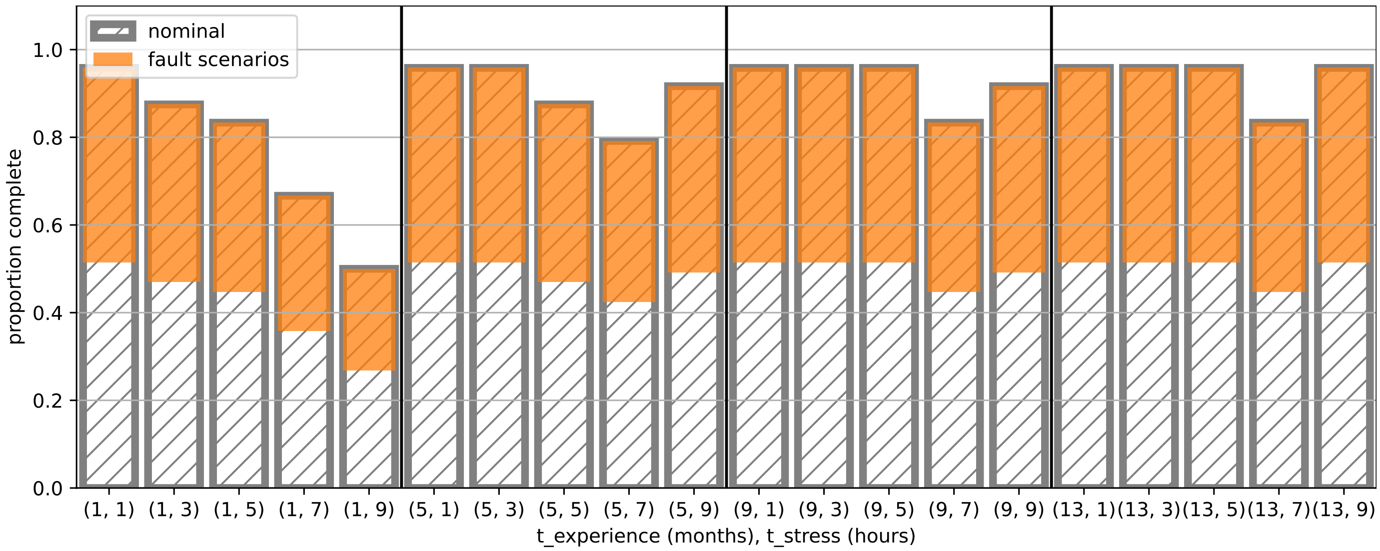

Fig. 9: Rover Resilience over Degraded Human Performance

Fig. 10: Nominal (no fault-injected) Missions over Human Degradation

oped over months) and operator stress over the course of a day effects the resulting rover resilience. The experience was modeled to accumulate over time through a learning curve (Sigmoid), where the rate of accumulation was driven by training frequency. Stress was modeled to increase exponentially over the course of the day. However, the rate of growth was driven by the experience (i.e., higher the experience, the slower the stress growth). To model this interaction between experience and stress, the short-term degradation of stress was nested within the model of long-term accumulation of experience, as shown in Figure 8. Experience, which accumulates over 25 months, was sampled at a few discrete times (month 1, 5, 8, and 13). The stress model was then simulated at these different levels of experience over the course of 10 hours. As shown in Figure 8b, the accumulation of stress (as modelled) is highly sensitive to the accumulation of experience, with large increases in stress over time corresponding to month 1 and 5 and much smaller increases occurring at the later months.

Joint times for stress and experience were then sampled and and simulated in the rover model in a nominal scenario, resulting in the results shown in Figure 10. As shown, mission failures are concentrated at the end of the day during the first few months of the operator's use of the rover. The resilience model was further simulated over the given drive faults, resulting in the responses in Figure 9. As shown in this figure, resilience degrades with the reliability of the system response, with the primary effect occurring over the course of the day in the first month of operations. This drop in resilience is roughly proportional to the drop in reliability, with a 51% of the fault scenarios at the first hour (with 96% reliability) still resulting in a mission completion and 27% of the fault scenarios still resulting in a mission completion at the ninth hour (with 50% reliability). This shows how degradation modelling can be used to evaluate human performance shaping factors which vary over time. While one might have differing assumptions about the various performance shaping factors involved (i.e., how quickly and how much stress will degrade performance or how stress accumulates overtime), this demonstration shows how the degradation (and accumulation) of human performance can be used to inform resilience assessment.

C. Combined Analysis

Finally, to demonstrate the joint consideration of human and component degradation at the same time, this section evaluates the rover at different levels of drive degradation along with stress and experience accumulation. This analysis thus occurs both across the different operator timescales for stress and experience accumulation and the component timescale for drive degradation. All three degradation models in Figure 6

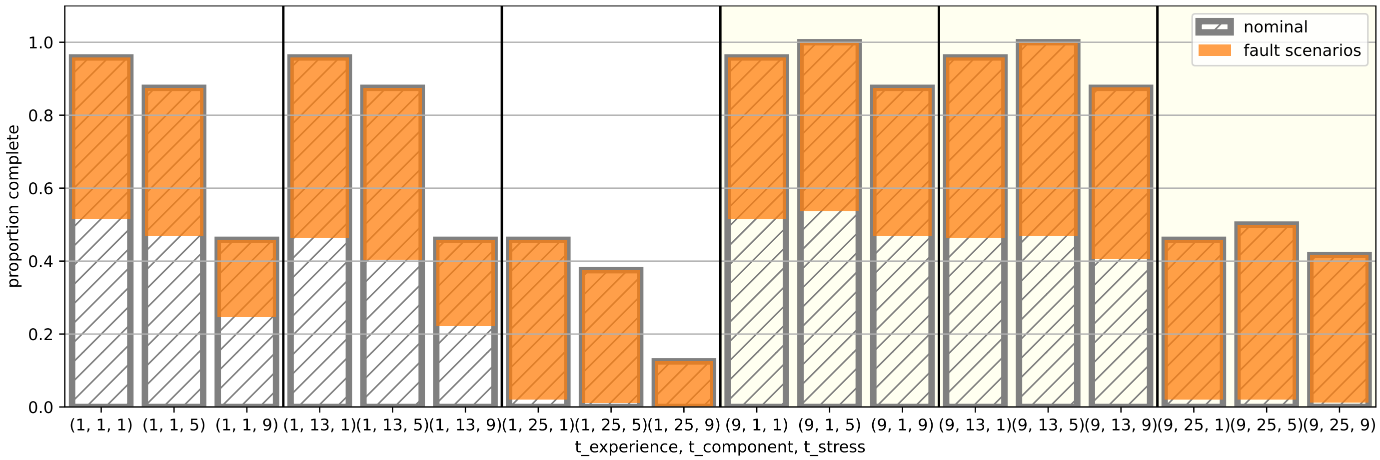

Fig. 11: Joint Effect of Component and Human Degradation on Resilience

and 8 were thus sampled at given times, and to reduce the computational complexity, the number of factor levels were reduced to two experience times (month 1 and 9), three component times (months 1,13, and 25), and three stress times (hour 1, 5, and 9). The results of this joint assessment are shown in Figure 11. As shown, the results are consistent with those in Figure 7 and Figure 9 with the drive system resulting in severe degradation at the 25^{th} month and human-caused degradation mostly occurring at the first month and ninth hour.

While the general trends for both human and component degradation behaviors are similar in the combined analysis, it should be noted that this joint consideration results in a more complete evaluation of resilience, since one is not assuming a perfect human in the component degradation analysis (or vice versa). As a result, once can see that the true resilience of the system is much lower than assumed when joint degradations are taken into account–for example, while the system may assumed to still be acceptably resilient at the 13^{th} month of component degradation (per Figure 7), this resilience is dependent on whether there are new or experienced operators. This is an important consideration–while it is often convenient to assess the impact of single parameters or single degradations on performance/resilience, this is an underestimate of hazards when multiple degradations (in multiple components or at multiple time-scales) could be occur at the same time. The resulting environment may be ripe for producing high-impact hazards where the (1) a critical component fails and (2) the human is unable to take the correct actions to mitigate the fault (like in the Air France 447 crash). Thus, evaluating the resilience of a system given system degradation in operations should (at some point) jointly consider the effect of all degradation behaviors to get an accurate assessment of operational resilience.

IV. Discussion

Considering degradation in resilience assessment models enables an important dimension of underlying performance variability to be taken into account–the degradation in system health over time from usage and operations. As shown in Section III, this consideration can reveal new hazards and resiliencies which would not have been accounted for by injecting fault scenarios in a "nominal" system. In addition to giving a more complete assessment of system resilience, this consideration can enable one to consider the effect of preventative strategies on system resilience. This component of resilience–preventative mitigation–is often considered a part of system resilience [28], [35], but is largely ignored by existing resilience frameworks, which instead focus on post-fault recovery. The advantage of considering the prevention-based component of resilience is that it enables more solutions to possible hazard-mitigating approaches to be considered in the design process, including:

1) Traditional operational strategies for managing degradation, such as prescribed maintenance and evaluation schedules for the physical system or training frequency and intensity, break requirements and shift length for the human system;
2) Traditional design strategies for reducing the impact of degradation, such as quality improvement, increase in design margin, and reduction of tolerances for the physical system and organizational and workspace design for the human system; and
3) Novel active design/operational strategies for managing degradation, such as Prognostics and Health Management for the physical system and human performance monitoring and warning systems for the human system.

While there are existing approaches for the consideration of these strategies on the basis of system reliability (in terms of prevention of faults), these approaches do not consider their effect on resilience (in terms of effect on hazard response), leading to an incomplete picture of degradation. Additionally, existing approaches often push the design of operational strategies (especially for human factors) late into the design phase, making it difficult to integrate these strategies with the design of the system. Considering these strategies in a high-level model like this ensures that these resilience strategies

integrate with the design of the system (rather than being considered as an afterthought). Finally, the consideration of novel, active strategies is key for considering resilience as an *adaptive* property of systems that not only recover from hazardous scenarios, but proactively mitigate them before they arise. In summary, degradation models enable a more comprehensive design for resilience process by enabling the consideration of both traditional passive and state-of-the art active design/operational strategies to proactively mitigate hazardous scenarios resulting from system degradation before they occur.

In the case of the Air France 447 crash considered in the Introduction, using a degradation-based resilience assessment could have helped the designers of the aircraft predict study the factors that eventually lead to the crash and account for them in the design and operations of the aircraft. Specifically, this framework would enable the designers to study the combined effects of degraded airspeed measurements and the experience, stress, and workload performance shaping factors that lead to the crash. Designers then could have uncovered that an inexperienced pilot is not capable of handling this situation, and set an operating requirement that more experienced pilots must in charge during severe weather or issue safety guidance and training on how to overcome such situations. They could also have developed diagnostic warning systems for the pitot probes that would warn pilots when weather conditions meant they were operating outside their range of accuracy. While there is no guarantee that the designers of the Air France 447 aircraft would have identified and mitigated this specific scenario using this framework (since we do not know how the designers would approach the model and analysis set-up), it would have enabled them to uncover events like it, resulting in an aircraft that is more resilient to similar types of failures.

V. Conclusions and Future Work

Previous approaches to resilience assessment have not considered the degradation in system performance that happens in operations due to known mechanisms of deterioration. To resolve this limitation, this paper proposed and demonstrated a framework for considering system performance degradation when assessing resilience in the early design phase. As demonstrated here on a model of a remotely-piloted rover, this help one understand how the degradation affects the modelled resilience. Depending on the types of degradation, this effect on degraded resilience could (1) have a substantial impact on safety margin and (2) precede the loss of reliability that often guides maintenance policy. Additionally, because degradation can decrease the system margin, and multiple degradations can interact with each other, neglecting operational degradation in resilience assessment can lead to an over-estimate of resilience (or resulting safety margin).

While the work presented here is an advancement on the methods in the literature, there are some limitations which should be acknowledged and addressed in future work. First, while was some demonstration of how to account for multiple scales of degradation in the analysis (see Figure 3), this assumed the degradation models had very little coupling between the simulations–future work should provide more guidance about how to manage this information when there are interactions between degradations occurring at different time-scales. Second, while this approach showed how to model degraded resilience over time, it did not provide a methodology for determining the overall expected resilience given this degradation. While this task may be simple enough (performing an expected value over the degradation time) it may become more complicated (and computationally costly) when there are multiple sources of faults and degradations and multiple design options must be compared against each other. Finally, future work should integrate this assessment with maintenance/workload scheduling to determine the optimal policy for system operations (e.g., as a part of a condition-based maintenance policy). This would enable its use both in design and operational phases to ensure that the desired level of resilience is maintained.

Acknowledgment

The U.S. government retains, and by accepting the article for publication, the publisher acknowledges that the U.S. government retains, a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this work, or allow others to do so, for U.S. government purposes.

References

[1] "Final report on the accident on 1st June 2009 to the Airbus A330-203 registered F-GZCP operated by Air France flight AF 447 Rio de Janeiro–Paris," Bureau d'Enquêtes et d'Analyses pour la sécurité de l'aviation civile, Paris: BEA, Investigation Report, 2012.

[2] N. Oliver, T. Calvard, and K. Potočnik, "Cognition, technology, and organizational limits: Lessons from the air france 447 disaster," *Organization Science*, vol. 28, no. 4, pp. 729–743, 2017.

[3] J. Wise, A. Rio, and M. Fedouach, "What really happened aboard air france 447," *Popular Mechanics*, vol. 6, 2011.

[4] M. Pecht, A. Dasgupta, D. Barker, and C. T. Leonard, "The reliability physics approach to failure prediction modelling," *Quality and Reliability Engineering International*, vol. 6, no. 4, pp. 267–273, 1990.

[5] M. Pecht and A. Dasgupta, "Physics-of-failure: an approach to reliable product development," *Journal of the IES*, vol. 38, no. 5, pp. 30–34, 1995.

[6] K. Ma, H. Wang, and F. Blaabjerg, "New approaches to reliability assessment: Using physics-of-failure for prediction and design in power electronics systems," *IEEE Power Electronics Magazine*, vol. 3, no. 4, pp. 28–41, 2016.

[7] H. Wang, M. Liserre, F. Blaabjerg, P. de Place Rimmen, J. B. Jacobsen, T. Kvisgaard, and J. Landkildehus, "Transitioning to physics-of-failure as a reliability driver in power electronics," *IEEE Journal of Emerging and Selected Topics in Power Electronics*, vol. 2, no. 1, pp. 97–114, 2013.

[8] I. Snook, J. M. Marshall, and R. Newman, "Physics of failure as an integrated part of design for reliability," in *Annual Reliability and Maintainability Symposium, 2003.* IEEE, 2003, pp. 46–54.

[9] M. Pecht and J. Gu, "Physics-of-failure-based prognostics for electronic products," *Transactions of the Institute of Measurement and Control*, vol. 31, no. 3-4, pp. 309–322, 2009.

[10] H. Oh, B. Han, P. McCluskey, C. Han, and B. D. Youn, "Physics-of-failure, condition monitoring, and prognostics of insulated gate bipolar transistor modules: A review," *IEEE Transactions on power electronics*, vol. 30, no. 5, pp. 2413–2426, 2014.

[11] P. Hall and J. Strutt, "Probabilistic physics-of-failure models for component reliabilities using monte carlo simulation and weibull analysis: a parametric study," *Reliability Engineering & System Safety*, vol. 80, no. 3, pp. 233–242, 2003.

[12] T. Tinga, "Predicting critical failures using physics of failure: opportunities and challenges," in *AVT-356 Research Symposium on Physics of Failure for Military Platform Critical Subsystems*. NATO Science & Technology Organization, 2021, pp. 1–13.

[13] N. Z. Gebraeel, M. A. Lawley, R. Li, and J. K. Ryan, "Residual-life distributions from component degradation signals: A bayesian approach," *IiE Transactions*, vol. 37, no. 6, pp. 543–557, 2005.

[14] Y.-F. Li, E. Zio, and Y.-H. Lin, "A multistate physics model of component degradation based on stochastic petri nets and simulation," *IEEE Transactions on Reliability*, vol. 61, no. 4, pp. 921–931, 2012.

[15] O. Fink, E. Zio, and U. Weidmann, "Predicting component reliability and level of degradation with complex-valued neural networks," *Reliability Engineering & System Safety*, vol. 121, pp. 198–206, 2014.

[16] N.-H. Kim, D. An, and J.-H. Choi, "Prognostics and health management of engineering systems," *Switzerland: Springer International Publishing*, 2017.

[17] C. Teubert, M. Corbetta, and C. Kulkarni, "Prognostics algorithm python package," May 2022. [Online]. Available: https://github.com/nasa/prog_algs

[18] A. Zwebek and P. Pilidis, "Degradation effects on combined cycle power plant performance—part i: gas turbine cycle component degradation effects," *J. Eng. Gas Turbines Power*, vol. 125, no. 3, pp. 651–657, 2003.

[19] R. Kurz, K. Brun, and M. Wollie, "Degradation effects on industrial gas turbines," *Journal of engineering for gas turbines and power*, vol. 131, no. 6, 2009.

[20] J. French, "A model to predict fatigue degraded performance," in *Proceedings of the IEEE 7th Conference on Human Factors and Power Plants*. IEEE, 2002, pp. 4–4.

[21] R. Patriarca, A. De Paolis, F. Costantino, and G. Di Gravio, "Simulation model for simple yet robust resilience assessment metrics for engineered systems," *Reliability Engineering & System Safety*, vol. 209, p. 107467, 2021.

[22] B. Cai, Y. Zhang, H. Wang, Y. Liu, R. Ji, C. Gao, X. Kong, and J. Liu, "Resilience evaluation methodology of engineering systems with dynamic-bayesian-network-based degradation and maintenance," *Reliability Engineering & System Safety*, vol. 209, p. 107464, 2021.

[23] Y. Cheng, E. A. Elsayed, and Z. Huang, "Systems resilience assessments: a review, framework and metrics," *International Journal of Production Research*, vol. 60, no. 2, pp. 595–622, 2022.

[24] A. Markina-Khusid, R. B. Jacobs, L. Antul, L. Cho, and H. T. Tran, "A complex network framework for validated assessments of systems of systems robustness," *IEEE Systems Journal*, 2021.

[25] S. Hosseini and D. Ivanov, "Bayesian networks for supply chain risk, resilience and ripple effect analysis: A literature review," *Expert systems with applications*, vol. 161, p. 113649, 2020.

[26] M. Bruneau, S. E. Chang, R. T. Eguchi, G. C. Lee, T. D. O'Rourke, A. M. Reinhorn, M. Shinozuka, K. Tierney, W. A. Wallace, and D. Von Winterfeldt, "A framework to quantitatively assess and enhance the seismic resilience of communities," *Earthquake spectra*, vol. 19, no. 4, pp. 733–752, 2003.

[27] N. Yodo and P. Wang, "Engineering resilience quantification and system design implications: A literature survey," *Journal of Mechanical Design*, vol. 138, no. 11, 2016.

[28] S. Hosseini, K. Barker, and J. E. Ramirez-Marquez, "A review of definitions and measures of system resilience," *Reliability Engineering & System Safety*, vol. 145, pp. 47–61, 2016.

[29] D. Hulse, H. Walsh, A. Dong, C. Hoyle, I. Tumer, C. Kulkarni, and K. Goebel, "fmdtools: A fault propagation toolkit for resilience assessment in early design," *International Journal of Prognostics and Health Management*, vol. 12, no. 3, 2021.

[30] S. Andrade and D. Hulse, "Evaluation and improvement of system-of-systems resilience in a simulation of wildfire emergency response," 2022, accepted Paper.

[31] D. Gertman, H. Blackman, J. Marble, J. Byers, C. Smith *et al.*, "The spar-h human reliability analysis method," *US Nuclear Regulatory Commission*, vol. 230, no. 4, p. 35, 2005.

[32] J. Liu, Y. Zou, W. Wang, L. Zhang, X. Liu, Q. Ding, Z. Qin, and M. Čepin, "Analysis of dependencies among performance shaping factors in human reliability analysis based on a system dynamics approach," *Reliability Engineering & System Safety*, vol. 215, p. 107890, 2021.

[33] M. De Ambroggi and P. Trucco, "Modelling and assessment of dependent performance shaping factors through analytic network process," *Reliability Engineering & System Safety*, vol. 96, no. 7, pp. 849–860, 2011.

[34] J. Park, W. Jung, and J. Kim, "Inter-relationships between performance shaping factors for human reliability analysis of nuclear power plants," *Nuclear Engineering and Technology*, vol. 52, no. 1, pp. 87–100, 2020.

[35] L. Shen, B. Cassottana, H. R. Heinimann, and L. C. Tang, "Large-scale systems resilience: A survey and unifying framework," *Quality and Reliability Engineering International*, vol. 36, no. 4, pp. 1386–1401, 2020.